

To: Dew, Judy[Dew.Judy@epa.gov]
From: Buch, Peter
Sent: Mon 1/30/2017 4:49:40 PM
Subject: Re: search report to show the highs of last week

Volume, response times and incidents

- Search volume on the primary Google Search Appliance rose from a normal business hours peak of 200 transactions/second to as high as 1,200 transactions per second in the Friday early hours. Transactions reached 1,000/second several times, 5 times normal peak volume.
- The escalation began on the morning of Tuesday, 1/24. The highest volume was recorded Thursday evening. Volume in general showed both business hours peaks and off-business hours peaks.

Incidents

- The Nagios monitors for the welch and shetland Linux servers tripped Tuesday morning and Tuesday afternoon due to slow response times. In all cases, search was responding, but a high percentage of responses were over the 5 second threshold (many were over 10 seconds).
- The Nagios monitors for welch and shetland tripped Friday afternoon due to a Java "too many files open" condition, cause be .
- The only actual outages were a "rolling" outage Tuesday evening to run of the Glossary search phase, and Friday afternoon to clear the too many files open condition.

Findings

- We discovered and blocked one very high volume metasearch bot from Amazon servers in Frankfurt, Germany. We discovered one low volume bot from DOE/OSTI which we did not block.
- The high off-hour activity suggests that there were other bots, but they were not clustered around single IPs. Not of the bots had the characteristics of denial of service attacks. Some of the queries could not have been submitted from the EPA search UI.
- There was extremely high organic activity in queries related to climate change.

Actions

- Turned off Dynamic Navigation, which increases latency on the Google Search Appliance
- BigIP changed the health check on the LTM that balances welch and shetland from a transaction to a "ping" that doesn't cause activity on the search appliance.
- Turned off Glossary search phase, which is one of 5 transactions to the appliance per search.

- Developed and tested a "Safe Mode" search that submits only on transaction to the appliance per search, but never reached the threshold where we needed to deploy it.

Peter Buch

Search Webmaster

CGI Federal |

Ex. 6 - Personal Privacy

(Work&Home)

Ex. 6 - Personal Privacy

(Cell)

Ex. 6 - Personal Privacy

buch.peter@epa.gov

From: Dew, Judy

Sent: Monday, January 30, 2017 8:39 AM

To: Buch, Peter

Subject: search report to show the highs of last week

Peter, I'm writing a weekly report for my managers and would like to include an item on the highs of search last week and what we had to do to keep search working.

One of the comments on Foresee last week was to improve search and I just laughed.

Judy Dew

Office of Information Management (OIM)

Web Content Services Division (WCSD)

Phone: (919) 541-2987

Fax: (919) 541-3648